ICE Buckets: Improved Counter Estimation for Network Measurement

Gil Einziger, Benny Fellman, Roy Friedman and Yaron Kassner

arXiv:1606.01364v1 [cs.NI] 4 Jun 2016

***Abstract*—Measurement capabilities are essential for a variety of network applications, such as load balancing, routing, fairness and intrusion detection. These capabilities require large counter arrays in order to monitor the traffic of all network flows. While commodity SRAM memories are capable of operating at line speed, they are too small to accommodate large counter arrays. Previous works suggested estimators, which trade precision for reduced space. However, in order to accurately estimate the largest counter, these methods compromise the accuracy of the smaller counters. In this work, we present a closed form representation of the optimal estimation function. We then introduce Independent Counter Estimation Buckets (ICE-Buckets), a novel algorithm that improves estimation accuracy for all counters. This is achieved by separating the flows to buckets and configuring the optimal estimation function according to each bucket's counter scale. We prove a tighter upper bound on the relative error and demonstrate an accuracy improvement of up to 57 times on real Internet packet traces.**

I. Introduction

A. Background

Counter arrays are essential in network measurements and accounting. Typically, measurement applications track several million flows [1], [2], and their counters are updated with the arrival of every packet. These capabilities are an important enabling factor for networking algorithms in many fields such as load balancing, routing, fairness, network caching and intrusion detection [3]–[7]. Counter arrays are also used in popular approximate counting sketches such as *multi stage filters* [8] and *count min sketch* [9], as well as in network monitoring architectures [10]–[12]. Such architectures are used to collect and analyze statistics from many networking devices [13].

Implementation of counter arrays is particularly challenging due to the requirement to operate at line speed. Although commodity SRAM memories are fast enough for this task, they do not meet the space requirements of modern counter arrays. Implementing a counter array entirely in SRAM is therefore very expensive [14].

Counter estimation algorithms use shorter counters, e.g., 12-bits instead of 32-bits, at the cost of a small error. Upon packet arrival, a counter is only incremented with a certain probability that depends on its current value. In order to keep the relative error uniform, small values are incremented with high probability and large ones with low probability. An *estimation function* is used in order to determine these probabilities and estimate the true value of a counter. Estimation functions can be *scaled* to achieve higher counting capacity at the cost of a larger estimation error.

Gil Einziger is with the department of Electrical Engineering, Politecnico di Torino, Italy. <gilga1983@gmail.com>.

Benny Fellman is with the department of Electrical Engineering, Technion, Israel. <benny.fellman@gmail.com>.

Roy Friedman and Yaron Kassner are with the department of computer science, Technion, Israel. <roy@cs.technion.ac.il>, <kassnery@gmail.com>

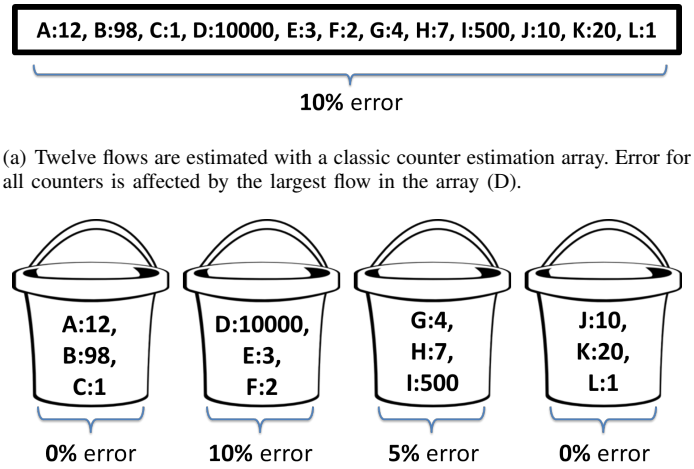

(a) Twelve flows are estimated with a classic counter estimation array. Error for all counters is affected by the largest flow in the array (D).

(b) The same twelve flows are estimated with ICE-Buckets. In this example the flows are separated into four buckets. Error for each counter is affected only by the largest flow in its bucket.

Fig. 1. An overview of ICE-Buckets vs. previous counter estimation approaches.

B. Contributions

In this work we present *Independent Counter Estimation Buckets (ICE-Buckets)*, a novel counter estimation technique that reduces the overall error by efficiently utilizing multiple counter scales.

The main principle of ICE-Buckets is illustrated in Figure 1. In this example, the largest counter (D) can only be estimated with a large scale and a relative error of 10%. In the traditional approach, this error applies to all counters, as illustrated in Figure 1(a). Figure 1(b) shows what happens when the array is partitioned into independent buckets. Counter D is still estimated with an error of 10%, but in this case the error applies only to counters within the same bucket. The other buckets are able to use smaller scales and enjoy lower relative error. Consequently, the overall error is reduced.

ICE-Buckets makes use of the optimal estimation function that was previously known only in recursive form. We present an explicit representation and provide an extended analysis for this function. We also present a rigorous mathematical analysis of ICE-Buckets that includes a very attractive upper bound for the overall relative error and a Chebyshev analysis to bound

the probability that the error is above a given threshold. We show that for traffic characteristics of real workloads this upper bound is up to 14 times smaller than that of previous works. Moreover, we show that the maximum relative error of ICE-Buckets is optimal.

We provide a lower and upper bounds for the space required to obtain a given counting capacity and error bound. We also analyze the error as a function of the maximal counting capacity and the number of estimation symbols. This analysis provides us with the mathematical tools to configure ICE-Buckets parameters in an optimal manner.

We further show how to perform decrements and down-scaling with ICE-buckets. Yet, their complicated mathematical analysis is left for future work.

Additionally, we extensively evaluate ICE-Buckets with five real Internet packet traces and demonstrate an accuracy improvement of up to 57 times. Finally, we show that ICE-Buckets can avoid global scale adjustments and still maintain similar accuracy. This configuration is more attractive for practical implementations.

In summary, we are the first to present a closed form explicit representation of an optimal estimation function. This enables us to extensively study the various aspects of this function using rigorous mathematical analysis, including the relation between its relative error, memory complexity, estimation symbol range, and even bound the probability of the actual error exceeding a certain value. We then propose the ICE-buckets technique, which divides counters into buckets, where each bucket is maintained with its own scale parameter, thereby greatly reducing the relative error. ICE-Buckets is also analyzed, and we show a methodological way of configuring its parameters. Finally, we simulate ICE-Buckets using 5 real-world traces and compare it to state of the art approaches, demonstrating its substantial benefits.

C. Related Work

While all counter arrays are required to monitor traffic at line speed, their implementations differ in the availability of monitored data. *Offline* counter arrays can take as much as several hours to read from, while *online* counter arrays can be read at line speed. Naturally, offline counter arrays are used for high level tasks such as data analysis and identifying performance bottlenecks. On the other hand, online counter arrays are used to answer low level queries such as what priority to give a certain flow, how much bandwidth it requires and where to route its packets.

Hybrid DRAM/SRAM counter arrays [1], [2] store only the least significant bits of each counter in SRAM and the rest of the counter in (slower) DRAM. In *CounterBraids* [15], counters are compressed in order to fit inside SRAM, but the decoding process is slow. Alternatively, *Randomized Counter Sharing (RCS)* [16] reduces the overhead required to maintain a flow to counter association. In that solution, each flow is randomly associated with a large number of counters and on each packet arrival a random counter is incremented. Statistical methods are then used in order to decode flow values. *Counter Tree* [17] further reduces the memory requirements of RCS by introducing the concept of *virtual counters*, each constructed from multiple physical counters organized in a tree structure such that large virtual counters span a path crossing multiple levels of the tree. Here, each flow is associated to multiple virtual counters using a plurality of hash functions. Hence, virtual counters share physical counters while flows share virtual counters and the virtual counters have variable size. Alas, CounterBraids, RCS, as well as Counter Tree, are all offline due to their long complex decode time, while hybrid SRAM/DRAM architectures are offline since reading requires accessing DRAM. Interestingly, estimators like the one suggested in this paper, can further improve the space efficiency of RCS and Counter Tree at the expense of precision.

Brick [18] is an online counter array that encodes variable length counters. Brick can hold more counters as the average counter is shorter than the largest one. Unfortunately, the counting capacity is limited and the encoding becomes less efficient as the average counter value increases. Alternatively, sampling techniques [12], [19] and heavy hitters algorithms [20]–[23] are able to monitor large flows. However, since they do not monitor all the flows, this type of solution is not always suitable.

Another popular approach for efficient flow statistics representation is shared counters. In these schemes, there is no longer a guaranteed one to one correspondence between a counter and a flow. Rather, some indirect hashing based mapping is maintained. This enables eliminating maintaining flow identifiers and the respective associations. Prominent example of these include *count min sketch* (CMS) [9], *multi-stage filters* [24], *spectral Bloom filters* (SBF) [25] and their variants, as well as TinyTable [26]. Shared counter techniques have a potential to work well with estimators, which can reduce the size of every shared counter.

Estimators are able to represent large values with small symbols at the price of a small error. They can therefore be used to implement online counter arrays. This idea was first introduced by *Approximate Counting* [27] and was recently adapted to networking as *Small Active Counters (SAC)* [28]. It was later improved by *DISCO* [29] in order to provide better accuracy and support variable sized increments.

[30] introduced a way to gradually increase the relative error as the counters grow. *CEDAR* [31] proved that their estimation function is optimal. *CASE* [32] extended our analysis of the optimal estimation function to also include variable increments. They showed that the large flows can be tracked by a cache to improve accuracy. CASE can be deployed with any estimation technique including the one presented in this paper.

In general, estimators require more space than sampling techniques, they provide accurate estimation for both small and large flows, and compared to Brick they enjoy significantly higher counting capacity at the price of a small relative error.

D. Paper Organization

The optimal estimation function is presented and analyzed in Section II, followed by the presentation and analysis of ICE-Buckets in Section III. Section IV describes simulation

results with real Internet packet traces. We conclude our work in Section V.

II. Optimal Estimation Function

Notation	Description
M	Maximum possible number of packets.
L	Number of different possible symbols - a power of two.
ε_{max}	Maximal relative error (for a single counter).
δ_{max}	Maximal coefficient of variation (CV) of the hitting time.
$A_\epsilon(l)$	The optimal estimation function with a scale of $\varepsilon = \varepsilon_{max}$

TABLE I
Notations

A. Technical Background

Consider the problem of counting up to M packets, with a counter of only $\log_2 L$ bits, where $\log_2 L < \log_2 (M+1)$ bits. We rely on an *estimation function* $A : \{0, ..., L-1\} \to [0, M]$, which accepts a symbol l as input and returns an *estimation value* for that symbol. M is the required *counting capacity* of the estimation function. For easy reference, the notations used in this section are summarized in Table I

First, the symbol l is initialized to zero. Upon arrival of a packet, we increment l with probability $\frac{1}{D(l)}$, where $D(l) = A(l+1) - A(l)$. It is easy to verify that the expected estimation value of l grows by one with each packet. Thus, the counter estimation is unbiased.

For example, for the estimation function $A(l) = 2^l$, if at a certain point in time a symbol $l = 3$ is used, its estimation value is $A(3) = 8$. If another packet arrives at the flow, we increment the symbol with probability $\frac{1}{D(3)} = \frac{1}{A(4)-A(3)} = \frac{1}{8}$. The estimation value is expected to change from 8 to 16 after 8 packet arrivals.

B. Our Estimation Function

We propose the following estimation function

$$A_\epsilon(l) = \frac{(1+2\epsilon^2)^l - 1}{2\epsilon^2}(1+\epsilon^2), \tag{1}$$

where ϵ is a parameter of the algorithm. The motivation and benefits of this estimation function are discussed in Subsection II-D.

C. Upscale

Upscale is a way to dynamically adjust the counter scale to the actual workload [30]. It is useful in case the counting capacity M is unknown. We begin with a small ϵ that gives a counting capacity of $A_\epsilon(L-1)$ and dynamically increase it when necessary. That is, when a symbol approaches $L-1$, we increase ϵ to $\epsilon' > \epsilon$. Then, we update all symbols to maintain unbiased estimation under the new scale.

Define l' to be the largest integer such that $A_{\epsilon'}(l') \le A_\epsilon(l)$. For our estimation function, this value is

$$l' = \left\lfloor log_{1+2\epsilon'^2}\left(1 + \frac{2\epsilon'^2 A_\epsilon(l)}{1+\epsilon'^2}\right)\right\rfloor \tag{2}$$

The correct estimation for symbol l lies between $A_{\epsilon'}(l')$ and $A_{\epsilon'}(l'+1)$. We update to $l'+1$ with probability proportional to the difference between $A_\epsilon(l)$ and $A_{\epsilon'}(l')$:

$$\frac{A_\epsilon(l) - A_{\epsilon'}(l')}{A_{\epsilon'}(l'+1) - A_{\epsilon'}(l')}$$

and to l' otherwise. Algorithm 1 describes the symbol upscale procedure.

Algorithm 1 Symbol Upscale

```
procedure SymbolUpscale(l,ε,ε')
    l' ← ⌊log_{1+2ε'^2}(1 + 2ε'^2 A_ε(l) / (1+ε'^2))⌋
    r ← rand(0,1)
    if r < (A_ε(l) − A_ε'(l')) / (A_ε'(l'+1) − A_ε'(l')) then
        l ← l' + 1
    else
        l ← l'
    end if
end procedure
```

D. Analysis

1) Performance Metrics: There are several metrics for the accuracy of an estimation function. Throughout this paper, we discuss the quality of estimation mainly in terms of the *root mean squared relative error (RMSRE)*, or *relative error* in short. Denote $\hat{n}$ the random variable representing the estimation value of a flow after n packets have arrived. The mean square relative error (MSRE) of a flow of size n is:

$$MSRE[n] = \mathbb{E}\left[\left(\frac{\hat{n}-n}{n}\right)^2\right]$$

and the root mean square relative error (RMSRE) is

$$RMSRE[n] = \sqrt{\mathbb{E}\left[\left(\frac{\hat{n}-n}{n}\right)^2\right]}$$

We want the *maximum relative error*,

$$\epsilon_{max} = \max_{n \le M} RMSRE[n],$$

to be as small as possible.

When counting multiple flows, we can also measure the *overall relative error*. Let n_i be the true value of counter i. The overall relative error is the root mean square relative error over all N counters,

$$\epsilon_{overall} = \sqrt{\frac{1}{N}\sum_i MSRE[n_i]}$$

Another metric for the accuracy of an estimation function is the *hitting time*. The *hitting time* is defined to be the random variable $T(l)$ that represents the amount of traffic required for a certain counter to be estimated as $A(l)$. The expected hitting time for symbol l is simply $A(l)$ in our case, according

to Theorem 2 in [31], and the *Coefficient of Variation (CV)* of the hitting time is

$$CV\left[T\left(l\right)\right]=\frac{\sigma[T(l)]}{\mathbb{E}[T(l)]}.$$

Simplistically, $CV\left[T\left(l\right)\right]$ measures the relative error when the symbol becomes l. In Theorem 1 below, we prove that our estimation function is optimal in terms of the *maximum CV of the hitting time*,

$$\delta_{max}=\max_{l<L}CV\left[T\left(l\right)\right].$$

2) Optimality: An estimation function is considered optimal if it minimizes δ_{max} given M, the desired counting capacity. To show that our estimation function is optimal, we rely on Theorems 3 and 4 in [31], stating that a function that satisfies the following recursive formula is optimal with $\delta_{max}=\delta$.

$$\forall l, A(l+1)-A(l)=\frac{1+2\delta^2A(l)}{1-\delta^2} \tag{3}$$

$$A\left(0\right)=0 \tag{4}$$

Furthermore, this function is unique.

Theorem 1. *$A_\epsilon(l)=\frac{(1+2\epsilon^2)^l-1}{2\epsilon^2}(1+\epsilon^2)$ is an optimal estimation function.*

Proof: Clearly, $A_\epsilon(0)=0$ and thus condition (4) holds.

$$\begin{aligned}A_\epsilon(l+1)&=\frac{\left((1+2\epsilon^2)^{l+1}-1\right)(1+\epsilon^2)}{2\epsilon^2}\\&=\frac{\left((1+2\epsilon^2)(1+2\epsilon^2)^l-(1+2\epsilon^2)+2\epsilon^2\right)(1+\epsilon^2)}{2\epsilon^2}\\&=(1+2\epsilon^2)A_\epsilon\left(l\right)+(1+\epsilon^2)\end{aligned} \tag{5}$$

Therefore,

$$A_\epsilon(l+1)-A_\epsilon(l)=2\epsilon^2A_\epsilon\left(l\right)+(1+\epsilon^2) \tag{6}$$

Choose

$$\delta=\sqrt{\frac{\epsilon^2}{1+\epsilon^2}}$$

to get

$$\epsilon^2=\frac{\delta^2}{1-\delta^2}.$$

Putting this into (6) we obtain

$$A_\epsilon(l+1)-A_\epsilon(l)=2\frac{\delta^2A_\epsilon\left(l\right)}{1-\delta^2}+1+\frac{\delta^2}{1-\delta^2}=\frac{1+2\delta^2A_\epsilon(l)}{1-\delta^2}.$$

Hence, condition (3) holds. Our estimation function satisfies conditions (3) and (4) and is therefore optimal. ∎

Thus, our estimation function is identical to the one given in CEDAR [31], which was previously known only in recursive form and was only analyzed with respect to its hitting time. In this work, we also analyze the relative error $RMSRE$, which is a natural metric to discuss.

3) Relative Error:

Theorem 2. *The optimal estimation function (A_ϵ) gives a relative error of*

$$RMSRE\left[n\right]=\epsilon,\qquad\forall n$$

Proof: To prove this theorem, we use a technique similar to the one used in [33].
Let $Q_l(n)$ be the probability to have a symbol l given that exactly n packets have arrived at the flow. As mentioned before, the estimator is unbiased, thus

$$\sum_l Q_l(n)A_\epsilon(l)=n. \tag{7}$$

In order to calculate ϵ, we should first find the variance of the estimation value. We already know its mean, so let us find

$$\mathbb{E}\left[\hat{n}^2\right]=\sum_l Q_l(n)A_\epsilon^2(l).$$

We first compute $\mathbb{E}\left[\left(n\hat{+}1\right)^2-\left(\hat{n}\right)^2\right]$. Recall that if the symbol is l, when a packet arrives the symbol is incremented with probability $\frac{1}{D_\epsilon(l)}=\frac{1}{A_\epsilon(l+1)-A_\epsilon(l)}$ or remains unchanged with probability $1-\frac{1}{D_\epsilon(l)}$. Therefore,

$$\begin{aligned}&\mathbb{E}\left[\left(n\hat{+}1\right)^2-\left(\hat{n}\right)^2\right]\\&=\sum_l\left(A_\epsilon(l+1)^2-A_\epsilon(l)^2\right)\cdot\frac{1}{D_\epsilon(l)}Q_l(n)\\&=\sum_l\left(A_\epsilon(l+1)+A_\epsilon(l)\right)Q_l(n)\end{aligned}$$

We can use (5) to substitute $A_\epsilon(l+1)$ with

$$(1+2\epsilon^2)A_\epsilon\left(l\right)+(1+\epsilon^2)$$

and obtain

$$\mathbb{E}\left[\left(n\hat{+}1\right)^2-\left(\hat{n}\right)^2\right]=\sum_l(1+\epsilon^2)(2A_\epsilon(l)+1)Q_l(n).$$

This can be separated to a constant multiplied by the unbiased mean, and another constant times the sum of a probability vector. We get

$$\begin{aligned}&\mathbb{E}\left[\left(n\hat{+}1\right)^2-\left(\hat{n}\right)^2\right]\\&=(1+\epsilon^2)\left(2\sum_l A_\epsilon(l)Q_l(n)+\sum_l Q_l(n)\right)\\&=(1+\epsilon^2)(2n+1)\end{aligned}$$

Now, we can calculate $\mathbb{E}\left[\hat{n}^2\right]$

$$\begin{aligned}\mathbb{E}\left[\hat{n}^2\right]&=\sum_{i=0}^{n-1}\left(\mathbb{E}[\left(i\hat{+}1\right)^2]-\mathbb{E}[\left(\hat{i}\right)^2]\right)+\mathbb{E}[\hat{0}^2]\\&=\sum_{i=0}^{n-1}(1+\epsilon^2)(2i+1)=(1+\epsilon^2)\frac{n\left(2n-1+1\right)}{2}\\&=(1+\epsilon^2)n^2\end{aligned}$$

Hence, the variance is

$$\mathbb{V}[\hat{n}] = \mathbb{E}\left[\hat{n}^2\right] - \mathbb{E}[\hat{n}]^2 = \epsilon^2 n^2$$

and the relative error is

$$RMSRE[n] = \sqrt{\frac{\mathbb{V}[\hat{n}]}{n^2}} = \epsilon$$

■

Note that the relative error is independent of n, and therefore $\epsilon_{max} = \epsilon$. For comparison, DISCO's [29] estimation function,

$$DISCO(l) = \frac{(1+2\epsilon^2)^l - 1}{2\epsilon^2},$$

also achieves a relative error bounded by ϵ. Thus, DISCO guarantees the same accuracy as the optimal estimation function but its counting capacity is $1+\epsilon^2$ times smaller. In most cases, this factor is negligible.

4) Upscale Error: We now show how to use linear programming to prove, for some ϵ and ϵ', that no upscale operation from ϵ to ϵ' increases the relative error to more than ϵ'.

Consider the change in variance when an upscale from ϵ to ϵ' occurs. Let Q_l denote the probability to have a symbol l before upscale. Recall that the probability to use l' as defined in Algorithm 1 is

$$\frac{A_{\epsilon'}(l'+1) - A_\epsilon(l)}{D_{\epsilon'}(l')}$$

and the probability to use $l'+1$ is

$$\frac{A_\epsilon(l) - A_{\epsilon'}(l')}{D_{\epsilon'}(l')}.$$

The expected estimation value remains unchanged after upscale [31]. Therefore, the change in variance is:

$$\begin{aligned}
\Delta\mathbb{V} &= \sum_l Q_l \frac{A_{\epsilon'}(l'+1) - A_\epsilon(l)}{D_{\epsilon'}(l')} \cdot \left(A^2_{\epsilon'}(l') - A^2_\epsilon(l)\right) \\
&\quad + Q_l \frac{A_\epsilon(l) - A_{\epsilon'}(l')}{D_{\epsilon'}(l')} \cdot \left(A^2_{\epsilon'}(l'+1) - A^2_\epsilon(l)\right) \\
&= \sum_l Q_l \frac{\left(A_{\epsilon'}(l'+1) - A_\epsilon(l)\right)\left(A_\epsilon(l) - A_{\epsilon'}(l')\right)}{D_{\epsilon'}(l')} \cdot \\
&\quad \cdot \left(-A_{\epsilon'}(l') - A_\epsilon(l) + A_{\epsilon'}(l'+1) + A_\epsilon(l)\right) \\
&= \sum_l Q_l \left(A_{\epsilon'}(l'+1) - A_\epsilon(l)\right)\left(A_\epsilon(l) - A_{\epsilon'}(l')\right)
\end{aligned}$$

If we find $\alpha \geq 0$ and $\beta \geq 0$ such that $\forall 0 \leq l < L$,

$$\left(A_{\epsilon'}(l'+1) - A_\epsilon(l)\right)\left(A_\epsilon(l) - A_{\epsilon'}(l')\right) \leq \alpha A^2_\epsilon(l) + \beta A_\epsilon(l) \tag{8}$$

we may conclude that the MSRE is

$$\begin{aligned}
MSRE[n] &\leq \frac{n^2\epsilon^2 + \Delta\mathbb{V}}{n^2} \\
&\leq \frac{n^2\epsilon^2 + \sum_l Q_l\left(\alpha A^2_\epsilon(l) + \beta A_\epsilon(l)\right)}{n^2}.
\end{aligned}$$

Recall that

$$\sum_l Q_l A^2_\epsilon(l) = \mathbb{E}\left[A^2_\epsilon(l)\right] = \mathbb{V}[\hat{n}] + \left(\mathbb{E}[\hat{n}]\right)^2 \leq n^2(1+\epsilon^2).$$

Therefore, if 8 holds for every l,

$$\begin{aligned}
MSRE[n] &\leq \frac{n^2\epsilon^2 + \alpha n^2\left(1+\epsilon^2\right) + \beta n}{n^2} \\
&= \epsilon^2 + \alpha\left(1+\epsilon^2\right) + \frac{\beta}{n} \\
&\leq \epsilon^2 + \alpha\left(1+\epsilon^2\right) + \beta.
\end{aligned}$$

Define the following two variable linear problem:

$$Minimize \qquad \epsilon^2 + \alpha\left(1+\epsilon^2\right) + \beta$$

such that the constraint (8) holds for every $0 \leq l < L$.

We solved this simple LP for a wide range of parameters and found that the objective is minimized to ϵ'^2 in all of these cases. We therefore conjecture that the relative error is always bounded by ϵ' after an upscale operation from ϵ to ϵ'.

5) Memory Complexity: We now evaluate how many bits per symbol ($\lceil \log_2 L \rceil$) are needed to count to M with a relative error of ϵ.

Theorem 3. *The required number of bits per symbol is*

$$\begin{aligned}
&\lceil \log_2 L \rceil \\
&= \left\lceil \log_2\left(1 + \frac{\ln\left((2M+1)\epsilon^2+1\right) - \ln\left(1+\epsilon^2\right)}{\ln\left(1+2\epsilon^2\right)}\right)\right\rceil \\
&= \log_2 \ln\left((2M+1)\epsilon^2+1\right) + \log_2 \epsilon^{-2} + \Theta(1)
\end{aligned}$$

Proof: We begin with the estimation value for symbol $L-1$, according to Equation (1):

$$M = \frac{\left(1+2\epsilon^2\right)^{L-1} - 1}{2\epsilon^2} \cdot \left(1+\epsilon^2\right).$$

We solve for L:

$$\frac{2M\epsilon^2}{1+\epsilon^2} + 1 = \left(1+2\epsilon^2\right)^{L-1}$$

$$\begin{aligned}
L &= 1 + \frac{\ln\left(\frac{2M\epsilon^2}{1+\epsilon^2} + 1\right)}{\ln\left(1+2\epsilon^2\right)} \\
&= 1 + \frac{\ln\left((2M+1)\epsilon^2+1\right) - \ln\left(1+\epsilon^2\right)}{\ln\left(1+2\epsilon^2\right)}.
\end{aligned}$$

Thus,

$$\begin{aligned}
&\lceil \log_2 L \rceil \\
&= \left\lceil \log_2\left(1 + \frac{\ln\left((2M+1)\epsilon^2+1\right) - \ln\left(1+\epsilon^2\right)}{\ln\left(1+2\epsilon^2\right)}\right)\right\rceil.
\end{aligned}$$

Next, we bound L from both directions to get a simpler expression. We start from below. We use the inequality

$$\ln(1+x) < x \text{ for } x > 0 \tag{9}$$

to obtain

$$\begin{aligned}
L &\geq 1 + \frac{\ln\left((2M+1)\epsilon^2+1\right) - \epsilon^2}{2\epsilon^2} \\
&= \frac{1}{2} + \frac{\ln\left((2M+1)\epsilon^2+1\right)}{2\epsilon^2}.
\end{aligned} \tag{10}$$

Next, we bound L from above. We use the inequality

$$\ln(1+x) > \frac{x}{1+x} \tag{11}$$

for $x > 0$,

$$\begin{aligned} L &\leq 1 + \frac{\ln\left((2M+1)\epsilon^2+1\right) - \frac{\epsilon^2}{1+\epsilon^2}}{\frac{2\epsilon^2}{1+2\epsilon^2}} \\ &\leq 1 + \frac{1+2\epsilon^2}{2\epsilon^2}\ln\left((2M+1)\epsilon^2+1\right) - \frac{1+2\epsilon^2}{2+2\epsilon^2} \end{aligned} \tag{12}$$

Notice that $\frac{1+2\epsilon^2}{2+2\epsilon^2} > \frac{1}{2}$, therefore we continue by replacing $\frac{1+2\epsilon^2}{2+2\epsilon^2}$ with $\frac{1}{2}$ in Inequality (12). We get:

$$L < \frac{1}{2} + \left(1+2\epsilon^2\right)\frac{\ln\left((2M+1)\epsilon^2+1\right)}{2\epsilon^2}. \tag{13}$$

Since $L \geq 2$, we get from Inequality (13) that

$$1.5 \leq \left(1+2\epsilon^2\right)\frac{\ln\left((2M+1)\epsilon^2+1\right)}{2\epsilon^2}.$$

Thus, replacing the number $\frac{1}{2}$ from Inequality (13) with the above expression divided by 3, we obtain

$$L \leq \frac{4\left(1+2\epsilon^2\right)}{6}\cdot\frac{\ln\left((2M+1)\epsilon^2+1\right)}{\epsilon^2}.$$

$\epsilon < 1$ and therefore

$$L < 2\cdot\frac{\ln\left((2M+1)\epsilon^2+1\right)}{\epsilon^2}.$$

Consequently, the amount of bits required to guarantee an error of at most ϵ and a counting capacity of at least M is

$$\lceil\log_2 L\rceil \leq \left\lceil 1 + \log_2\ln\left((2M+1)\epsilon^2+1\right) + \log_2\epsilon^{-2}\right\rceil \tag{14}$$

and on the other hand (from Inequality (10))

$$\lceil\log_2 L\rceil \geq \left\lceil -1 + \log_2\ln\left((2M+1)\epsilon^2+1\right) + \log_2\epsilon^{-2}\right\rceil. \tag{15}$$

Therefore, the amount of memory required is

$$\lceil\log_2 L\rceil = \log_2\ln\left((2M+1)\epsilon^2+1\right) + \log_2\epsilon^{-2} + \Theta(1).$$

■

Corollary 1. *For a constant ϵ, the required number of bits is $O(\log\log M)$.*

In figure 2, we demonstrate the accuracy of the memory bounds we derive in Theorem 3 (inequalities (14) and (15)). We simulate different numbers of bits per symbol with $\epsilon = 2^{-5}$. For each value of L we then calculate the counting capacity M and from it the memory bounds from inequalities (14) and (15). For these parameters, we can see that the lower bound is tight, and the upper bound is only two bits larger. In addition, the figure shows that with an approximation of $\epsilon = 2^{-5}$, 13 bits are sufficient to count up to more than 2^{32}.

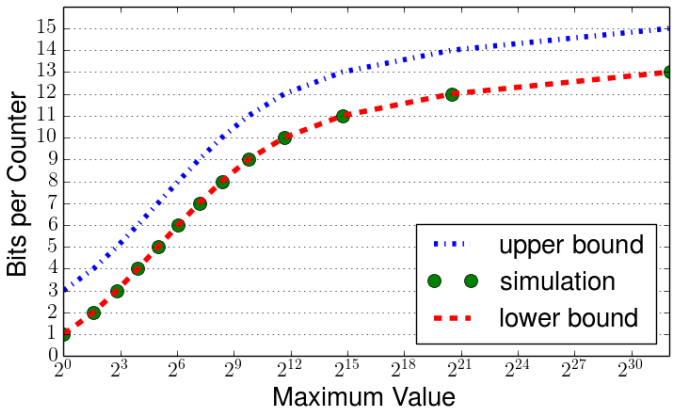

Fig. 2. Simulation and bounds for the number of bits per symbol required to count with different values of M for $\epsilon = 2^{-5}$.

6) ϵ as a function of M and L: Computing the smallest ϵ sufficient to represent a counter M with L possible symbols may be useful for several purposes. First, if we know M in advance and we have limited memory, we can detect the optimal parameter to use and avoid the upscale phase. Second, it enables theoretical comparison of the algorithm's relative error.

To achieve a counting capacity M with $\log_2 L$ bits, the error should be the ϵ that solves

$$M = \frac{\left(1+2\epsilon^2\right)^{L-1}-1}{2\epsilon^2}\cdot\left(1+\epsilon^2\right). \tag{16}$$

$$\frac{2M\epsilon^2}{1+\epsilon^2} = \left(1+2\epsilon^2\right)^{L-1} - 1.$$

We show two different ϵ's and show that one gives a value greater than M when put into Equation (16) and the other gives a value smaller than M.

Claim 1.

$$\epsilon^2 \leq 3\frac{\ln\frac{6M}{L-1}}{L-1}.$$

Proof: First, let

$$\epsilon^2 = 3\frac{\ln\frac{6M}{L-1}}{L-1}.$$

By developing Equation (16), we get

$$\begin{aligned} A_\epsilon(L-1) &= \frac{\left(1+2\epsilon^2\right)^{L-1}-1}{2\epsilon^2}\cdot\left(1+\epsilon^2\right) \\ &= \frac{e^{\ln\left(1+2\epsilon^2\right)\cdot(L-1)}-1}{2\epsilon^2}\cdot\left(1+\epsilon^2\right). \end{aligned}$$

Now we use Inequality (11) to obtain

$$A_\epsilon(L-1) > \frac{e^{\frac{2\epsilon^2}{1+2\epsilon^2}\cdot(L-1)}-1}{2\epsilon^2}.$$

Since $\epsilon < 1$, we have

$$A_\epsilon(L-1) > \frac{e^{\frac{2\epsilon^2}{3}\cdot(L-1)}-1}{2\epsilon^2} = \frac{e^{2\ln\frac{6M}{L-1}}-1}{6\frac{\ln\frac{6M}{L-1}}{L-1}} = \frac{\left(\left(\frac{6M}{L-1}\right)^2-1\right)(L-1)}{6\ln\frac{6M}{L-1}}.$$

Now, we can use Inequality (9) to obtain

$$A_\epsilon(L-1) > \frac{\left(\frac{6M}{L-1}-1\right)\left(\frac{6M}{L-1}+1\right)(L-1)}{6\left(\frac{6M}{L-1}-1\right)} > M.$$

Since M is increasing with ϵ and the chosen ϵ gives a counting capacity greater than M, to achieve a counting capacity of exactly M, ϵ must be less than or equal to $3\frac{\ln\frac{6M}{L-1}}{L-1}$. ■

Claim 2.

$$\epsilon^2 \geq \frac{\ln\frac{2M+1}{2L-1}}{2(L-1)}.$$

Proof: Let

$$\epsilon^2 = \frac{\ln\frac{2M+1}{2L-1}}{2(L-1)}.$$

As before,

$$A_\epsilon(L-1) = \frac{\left(1+2\epsilon^2\right)^{L-1}-1}{2\epsilon^2}\cdot\left(1+\epsilon^2\right) = \frac{e^{\ln\left(1+2\epsilon^2\right)\cdot(L-1)}-1}{2\epsilon^2}\cdot\left(1+\epsilon^2\right).$$

We use Inequality (9) to get

$$A_\epsilon(L-1) \leq \frac{e^{2\epsilon^2\cdot(L-1)}-1}{2\epsilon^2}\cdot\left(1+\epsilon^2\right) = \frac{e^{\ln\frac{2M+1}{2L-1}}-1}{2\frac{\ln\frac{2M+1}{2L-1}}{2(L-1)}}\cdot\left(1+\frac{\ln\frac{2M+1}{2L-1}}{2(L-1)}\right) = \frac{\frac{2M+1}{2L-1}-1}{2}\cdot\left(1+\frac{2(L-1)}{\ln\frac{2M+1}{2L-1}}\right)$$

Returning to Inequality (11), we can get

$$A_\epsilon(L-1) \leq \frac{\frac{2M+1}{2L-1}-1}{2}\cdot\left(1+\frac{2(L-1)\cdot\frac{2M+1}{2L-1}}{\frac{2M+1}{2L-1}-1}\right) = \frac{\frac{2M+1}{2L-1}-1+2(L-1)\cdot\frac{2M+1}{2L-1}}{2} = M.$$

We have shown that an ϵ of $\frac{\ln\frac{2M+1}{2L-1}}{2(L-1)}$ gives a counting capacity of at most M. Since M is increasing with ϵ, we conclude that ϵ must be greater or equal to $\frac{\ln\frac{2M+1}{2L-1}}{2(L-1)}$. ■

Using Bolzano-Weierstrass theorem, we conclude that there exists an ϵ such that

$$\frac{\ln\frac{2M+1}{2L-1}}{2(L-1)} \leq \epsilon^2 \leq 3\frac{\ln\frac{6M}{L-1}}{L-1}$$

for which Equation (16) holds.

The two sides of the inequality may be asymptotically distinct. But, for interesting cases, this does not happen. If $M \leq 2L-1$, then it is sufficient to use one more bit than $\lceil \log_2 L \rceil$ to represent M exactly. Therefore, the interesting case is when $M > 2L-1$. In this case, it is simple to see that both sides of the above equation are asymptotically $\Theta\left(\frac{\ln\frac{M}{L}}{L}\right)$. Hence, for $M > 2L-1$,

$$\epsilon^2 = \Theta\left(\frac{\ln\frac{M}{L}}{L}\right).$$

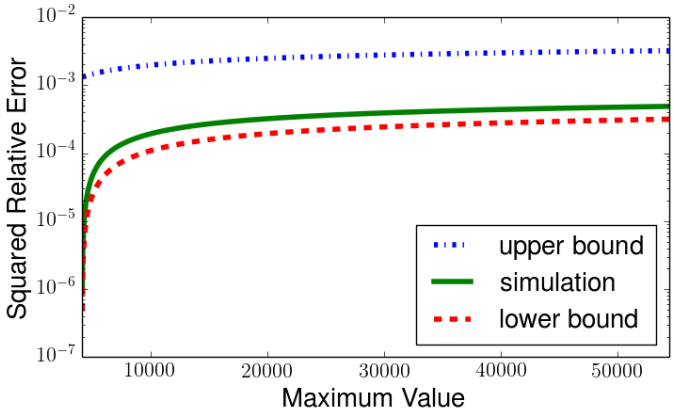

Fig. 3. Squared relative error and bounds with a 12-bit symbol for different values of M.

Figure 3 demonstrates the findings. To create it, we simulated M with $L = 4096$ and errors ranging between 2^{-20} and 2^{-11}. For each M, we then calculated the lower and upper bound. We first see that the true error indeed lies between the lower and the upper bound. In this case, the lower bound is much tighter than the upper bound. For small values of M, where $M < 2L-1 = 9191$, the upper bound is inaccurate. However, for larger values of M, we see that the two bounds and the true error are asymptotically the same, as shown above.

7) Chebyshev Analysis: The RMSRE metric may be unsuitable for some applications because it only describes the maximum *expected* relative error rather than the maximum relative error. Some applications may require determining with certainty $1-\rho$ that the relative error $\frac{|\hat{n}-n|}{n}$ is no more than β. We can obtain that guarantee using Chebishev's inequality, according to which

$$Pr\left(\frac{|\hat{n}-n|}{n} \geq k\epsilon\right) = Pr\left(|\hat{n}-n| \geq k\sigma\right) \leq \frac{1}{k^2},$$

where $\sigma = n\epsilon$ is the standard deviation. Choosing $k = \frac{\beta}{\epsilon}$ gives us a probability $\rho = \frac{\epsilon^2}{\beta^2}$. This allows us to choose the parameter ϵ according to any given β and ρ. For example, to achieve a relative error of more than $\beta = 10\%$ with probability no more than $\rho = 1\%$, we can use $\epsilon = \sqrt{\beta^2\rho} = 1\%$.

E. Decrementing Counters

Decrementing counters can be useful for applications that need to "forget" old, less relevant, values, e.g., when we wish to count flows in a sliding window [23]. The process

of decrementing a flow is very similar to the process of incrementing a flow. Instead of incrementing l with probability $\frac{1}{D(l)}$, we decrement it with probability $\frac{1}{D(l-1)}$.

Theorem 4. *For an unbounded counter, the estimation is unbiased after any number of increments or decrements.*

Proof: Let l_t be the random variable that represents the counter at time t. Let n_t be the number of increments minus the number of decrements until and including time t. Assume by induction on t that the estimation is unbiased at time t, i.e. $\mathbb{E}\left[A(l_t)\right] = n_t$. We next show that the estimation remains unbiased at time $t+1$. If the symbol is incremented at time $t+1$,

$$\begin{aligned}
\mathbb{E}\left[A(l_{t+1})\right] &= \sum_x A(l_{t+1}) Pr\left(l_t = x\right) \\
&= \sum_x \left(\frac{1}{D(x)} A(x+1) + \left(1 - \frac{1}{D(x)}\right) A(x)\right) Pr\left(l_t = x\right) \\
&= \sum_x \frac{1}{D(x)} \left(A(x+1) - A(x)\right) Pr\left(l_t = x\right) \\
&\quad + \sum_x A(x) Pr\left(l_t = x\right) \\
&= \sum_x Pr\left(l_t = x\right) + n_t = 1 + n_t
\end{aligned}$$

Similarly, if the symbol is decremented at time $t+1$,

$$\begin{aligned}
\mathbb{E}\left[A(l_{t+1})\right] &= \sum_x A(l_{t+1}) Pr\left(l_t = x\right) \\
&= \sum_x \frac{1}{D(x-1)} A(x-1) Pr\left(l_t = x\right) \\
&\quad + \sum_x \left(1 - \frac{1}{D(x-1)}\right) A(x) Pr\left(l_t = x\right) \\
&= \sum_x \frac{1}{D(x-1)} \left(A(x-1) - A(x)\right) Pr\left(l_t = x\right) \\
&\quad + \sum_x A(x) Pr\left(l_t = x\right) \\
&= \sum_x -Pr\left(l_t = x\right) + n_t = -1 + n_t
\end{aligned}$$

Therefore, the estimation is unbiased at time $t+1$. At time 0 the estimation is 0 and therefore unbiased. In conclusion, the estimation is unbiased at every time t. ∎

F. Downscale

In scenarios where counters are decremented, we may find ourselves in a situation where counters are small yet their scale is large, resulting in a quickly increasing error. In this case, we may want to *downscale* the counters.

Let ϵ be the current error parameter and ϵ' a smaller, desired error. We downscale all counters after checking that they can all be represented with $A_{\epsilon'}$. The process of checking whether all counters can be represented with $A_{\epsilon'}$ is repeated until all counters are small enough. In this process, we iterate over every symbol l and check that $A_\epsilon(l) < A_{\epsilon'}(L-1)$. If this iteration and downscaling of the counters takes time and U updates are performed during that time, we might want to check instead that $A_\epsilon(l+U) < A_{\epsilon'}(L-1)$. This guarantees that once all counters are downscaled, they can still be represented with $A_{\epsilon'}$.

Then, we update all symbols to maintain unbiased estimation under the new scale. This is done with the Symbol Downscale procedure, which is identical to the Symbol Upscale procedure defined in Algorithm 1. The only difference is that this time $\epsilon' < \epsilon$.

However, changing the parameter of the estimation function to ϵ' does not reduce the relative error to ϵ'. It is important to distinct between the error parameter and the actual error.

Reducing the scale of a counter complicates the analysis. Therefore, in this work we assume that no decrements and no downscale occur, i.e., packets only arrive and do not leave.

III. ICE-Buckets

A. Overview

We now describe ICE-Buckets, a data structure that uses the optimal estimation function with a different scale for each bucket. First, notations specific to this section are given in Table II. ICE-Buckets uses a base error parameter that is called ϵ_{step}. Symbols are separated into buckets and each bucket maintains a scale parameter w_i of size $\log_2 E$ bits, where E is a parameter of the data structure. To estimate counters in a bucket with scale w, we use the estimation function $A_{\epsilon_{step} \cdot w}$. We pay special attention to the additional memory that is allocated for each bucket and make sure that this overhead is small. Figure 4 demonstrates this basic architecture.

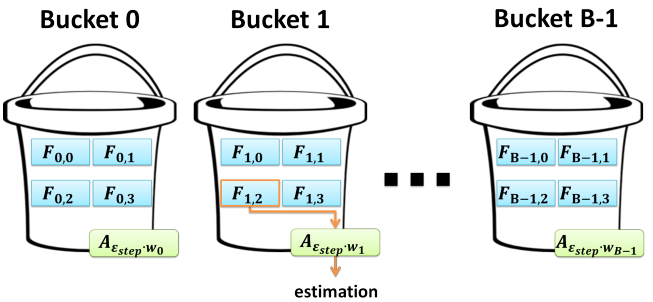

Fig. 4. An ICE-Buckets data structure with four flows per bucket. Bucket i has a scale parameter w_i, which is used by the estimation function $A_{\epsilon_{step} \cdot w_i}$ to decode the symbols $(F_{i,j})$. In this example, to estimate counter 2 in bucket 1, $A_{\epsilon_{step} \cdot w_1}(F_{1,2})$ is computed.

B. Algorithm

Initially, all symbols are set to zero. Since each bucket's scale is different, we first associate each flow-id with a bucket. To estimate the value of flow f, we use $A_{\epsilon_{w_i}}(F_{ij})$ where $i = \left\lfloor \frac{f}{S} \right\rfloor$ and $j = (f \mod S)$.

The estimation values for bucket i are calculated with an optimal estimation function of scale $\epsilon_{w_i} = \epsilon_{step} \cdot w_i$. An optimal function with $\epsilon = 0$ is defined as the identity function, i.e., $\forall l : A_0(l) = l$.

Notation	Description
M	Maximum possible number of packets.
L	Number of different possible symbols - a power of two.
S	Number of symbols per bucket.
N	Number of flows.
B	Number of buckets $\left(\frac{N}{S}\right)$. For simplicity assume that N is a multitude of S.
F_{ij}	Symbol array of size $B \times S$. i runs over the buckets and j runs over the flows in each bucket. Each symbol is $\log_2 L$ bits wide.
w_i	Scale parameter for bucket i.
ϵ_{step}	The difference between consecutive estimation errors.
ϵ_w	Estimation error for bucket with scale parameter w ($\epsilon_{step} \cdot w$).
$\varepsilon_{overall}$	Overall relative error (averaged on all counters).
E	Number of different possible estimation scales - a power of two.
T	Space allocated for the data structure (bits).

TABLE II
NOTATIONS

The increment process is straightforward. When a packet arrives, first we find the associated bucket i and the index in the bucket, j. Then, we increment F_{ij} with probability $\frac{1}{D_{\epsilon_{w_i}}(F_{ij})}$.

C. Analysis

In this section, we analyze ICE-Buckets and bound both its maximum relative error and overall relative error.

1) Maximum Relative Error: Define the following function,

$$m(\epsilon) = A_\epsilon(L-1) = \frac{(1+2\epsilon^2)^{L-1}-1}{2\epsilon^2}(1+\epsilon^2),$$

which denotes the maximum representable value with error ϵ.

Define $\varepsilon(M)$ to be the smallest ϵ we need for an optimal estimation function with capacity at least M. For $M > L-1$, $\varepsilon(M)$ is the inverse of $m(\epsilon)$. For $M \leq L-1$, $\varepsilon(M)$ is zero. $m(\epsilon)$ is increasing with ϵ and therefore $\varepsilon(M)$ is increasing with M. We now show that ICE-Buckets' relative error is not larger than $\varepsilon(M)$.

Theorem 5. *The maximum relative error of ICE-Buckets can be bounded by $\varepsilon(M)$ for any distribution of the counters.*

Proof: Construct an ICE-Buckets structure with $\epsilon_{step} = \frac{\varepsilon(M)}{E-1}$, where E is the number of different possible estimation scales. Let M_i be the value of the biggest counter in bucket i. If we choose the scale of bucket i to be at least $w_i = \left\lceil \frac{\varepsilon(M_i)}{\epsilon_{step}} \right\rceil$, we obtain an error parameter of no less than

$$\left\lceil \frac{\epsilon(M_i)}{\epsilon_{step}} \right\rceil \epsilon_{step} \geq \epsilon(M_i),$$

which gives us a capacity of at least M_i. In other words, we round up each bucket's error to the nearest product of ϵ_{step}.

$\varepsilon(M_i) \leq \varepsilon(M)$, because the maximum counter in each bucket is smaller or equal to the total number of packets and ϵ is increasing with M. Thus, the maximum relative error over the entire counter scale is bounded by $\varepsilon(M)$. ∎

We conclude that ICE-Buckets has the same maximum relative error as the optimal estimation function.

2) Overall Relative Error: ICE-Buckets also improves the guaranteed overall relative error. In ICE-Buckets, this error is

$$\epsilon_{overall} = \sqrt{\frac{1}{N}\sum_f MSRE[n_f]} \tag{17}$$

In Theorem 6, we show an upper bound for $\epsilon_{overall}$ when ICE-Buckets is configured optimally. In order to prove it, we need to show that $\epsilon^2(M)$ is concave for $M \geq L-1$ (as can be seen in Figure 5). We observe that $m(\epsilon)$ is convex and increasing with ϵ^2. The inverse of this function (which is defined on $M \geq L-1$) is therefore concave and increasing. Similarly, $\varepsilon(M)$ is also increasing and concave on $M \geq L-1$.

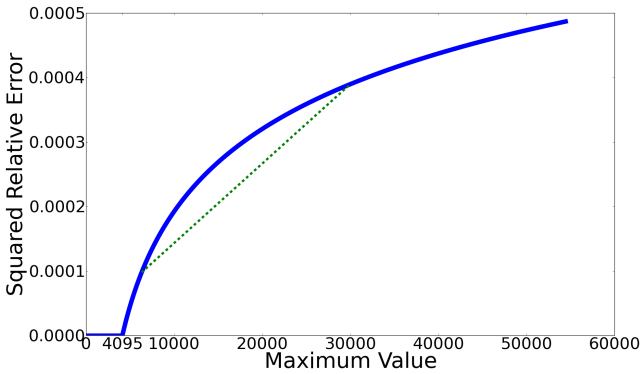

Fig. 5. ϵ^2 as a function of M for $L = 4096$. The dashed line comes to show that the function is concave for $M \geq L-1$.

Theorem 6. *For any counter distribution, ICE-Buckets can be configured to have an overall relative error no greater than*

$$\epsilon\left(\frac{M}{B}+L-1\right)+\frac{\varepsilon(M)}{E-1}.$$

Proof: Use the same construction as in the proof of Theorem 5, i.e., $\epsilon_{step} = \frac{\varepsilon(M)}{E-1}$ and $w_i = \left\lceil \frac{\varepsilon(M_i)}{\epsilon_{step}} \right\rceil$. With this construction, the overall relative error is no greater than

$$\epsilon_{overall} \leq \sqrt{\frac{1}{N}\sum_f MSRE[n_f]} \leq \sqrt{\frac{\sum_{i=0}^{B-1}\left\lceil \frac{\epsilon(M_i)}{\epsilon_{step}} \right\rceil^2 \epsilon_{step}^2}{B}}.$$

Instead of rounding $\varepsilon(M_i)$ to be a product of ϵ_{step}, we can

simply add ϵ_{step} and keep a bound that is no smaller.

$$\epsilon_{overall} \leq \sqrt{\frac{\sum_{i=0}^{B-1} \left\lceil \frac{\epsilon(M_i)}{\epsilon_{step}} \right\rceil^2 \epsilon_{step}^2}{B}} \leq \sqrt{\frac{\sum_{i=0}^{B-1} \left(\epsilon\left(M_i\right) + \epsilon_{step}\right)^2}{B}}$$

We claim that according to the concaveness of ϵ and ϵ^2, $\left(\varepsilon(x) + \epsilon_{step}\right)^2$ is concave on $x > L-1$. This is because:

$$\begin{aligned}
&\alpha\left(\varepsilon(x) + \epsilon_{step}\right)^2 + (1-\alpha)\left(\varepsilon(y) + \epsilon_{step}\right)^2 \\
=&\alpha\epsilon^2(x) + (1-\alpha)\,\epsilon^2(y) \\
&+ 2\epsilon_{step}\left(\alpha\varepsilon(x) + (1-\alpha)\,\varepsilon(y)\right) + \epsilon_{step}^2 \\
\leq&\epsilon^2\left(\alpha x + (1-\alpha)\,y\right) + 2\epsilon_{step}\varepsilon\left(\alpha x + (1-\alpha)\,y\right) + \epsilon_{step}^2 \\
\leq&\left(\varepsilon\left(\alpha x + (1-\alpha)\,y\right) + \epsilon_{step}\right)^2.
\end{aligned}$$

To use concaveness, we must make sure that all of the values are greater or equal to $L-1$. We do so by adding $L-1$ to each symbol. Since ϵ is an increasing function, the result is no smaller. That is,

$$\epsilon_{overall} \leq \sqrt{\frac{\sum_{i=0}^{B-1} \left(\epsilon\left(M_i + L - 1\right) + \epsilon_{step}\right)^2}{B}}.$$

We can now apply Jensen's inequality:

$$\begin{aligned}
\epsilon_{overall} \leq& \sqrt{\frac{\sum_{i=0}^{B-1} \left(\epsilon\left(M_i + L - 1\right) + \epsilon_{step}\right)^2}{B}} \\
\leq& \sqrt{\left(\epsilon\left(\frac{\sum_{i=0}^{B-1}\left(M_i + L - 1\right)}{B}\right) + \epsilon_{step}\right)^2} \\
=& \epsilon\left(\frac{M}{B} + L - 1\right) + \epsilon_{step}.
\end{aligned}$$

Setting $\epsilon_{step} = \frac{\varepsilon(M)}{E-1}$, we get an overall relative error of no more than

$$\epsilon_{overall} \leq \epsilon\left(\frac{M}{B} + L - 1\right) + \frac{\varepsilon\left(M\right)}{E-1}.$$

■

With the correct choice of parameters, this guaranteed overall relative error is far better than the one we can achieve with CEDAR - $\varepsilon(M)$.

This bound demonstrates the effect E and S have on the error. We want the bucket size $S = \frac{N}{B}$ to be as small as possible to restrict the negative impact on other counters, as a counter's size only affects the errors of counters sharing the same bucket. As for E, we want it to be as big as possible to increase the error granularity. However, decreasing S or increasing E increases the memory overhead.

Note that while CEDAR guarantees optimal estimation in terms of the maximum CV of the hitting time, there is no such claim in terms of the overall relative error. This enables ICE-Buckets to significantly improve $\epsilon_{overall}$.

D. Dynamic Configuration

In Theorem 6, we have shown that there is an ICE-Buckets configuration for which the error is low. We now show how to dynamically configure ICE-Buckets to fit any workload. This process is composed of local and global upscale operations.

1) Local Upscale: The configuration of the data structure, $\{w_i\}_{i=0}^{B-1}$, is dynamically adjusted to the biggest estimation value in each bucket. Initially, and bucket scales are set to Zero. Whenever a symbol F_{ij} approaches L, we increment w_i and upscale bucket i to use the parameter ϵ_{w_i+1}. This is done by upscaling all of the flows in bucket i using the symbol-upscale procedure described in Algorithm 1. A pseudo code of the local upscale procedure can be found in Algorithm 2. We note that since the number of counters per bucket S is small, local upscale can be efficiently implemented in hardware.

Algorithm 2 Local Upscale

```
1: procedure UpscaleBucket(i)
2:     for j = 0 to S − 1 do
3:         SymbolUpscale(F_ij, ε_{w_i}, ε_{w_i+1})
4:     end for
5:     w_i ← w_i + 1
6: end procedure
```

2) Global Upscale: When a counter in a bucket with the maximum scale index $(E-1)$ approaches its maximum value $(L-1)$, we initiate a global upscale procedure to prevent overflow. The procedure doubles the size of ϵ_{step}. Buckets with odd w_is perform a local upscale. Then, every bucket i updates its scale index to $w_i/2$. Pseudo code is given in Algorithm 3.

Algorithm 3 Global Upscale

```
1: procedure GlobalUpscale
2:     for u = 0 to B − 1 do
3:         if w_u mod 2 = 1 then
4:             UpscaleBucket(u)
5:         end if
6:         w_u ← w_u / 2
7:     end for
8:     ε_step ← 2ε_step
9: end procedure
```

Table III demonstrates this process. In this case, we use eight different possible scales ($E = 8$) and $\epsilon_{step} = 0.1\%$ before the global upscale. Buckets with odd w_i perform local upscale and their scale is incremented accordingly, so that all buckets have even scale parameters and we can safely half their scales to match the new ϵ_{step}. Notice that at the moment of upscale, ϵ_{max} is increased only by ϵ_{step}. In this case, no bucket after upscale has a scale parameter larger than 4. In general, no bucket has a scale parameter larger than $\frac{E}{2}$ immediately after global upscale.

Global upscale may be difficult to implement in hardware. A method to upscale the entire counter array while continuously counting new packet arrivals is described in [31]. This method also applies to ICE-Buckets' global upscale. In our case, global upscale can be completely avoided by setting the maximal error to $\varepsilon(M)$.

Old w	Old ϵ_w	New w	New ϵ_w	Requires Upscale?
0	0.0%	0	0.0%	No
1	0.1%	1	0.2%	Yes
2	0.2%	1	0.2%	No
3	0.3%	2	0.4%	Yes
4	0.4%	2	0.4%	No
5	0.5%	3	0.6%	Yes
6	0.6%	3	0.6%	No
7	0.7%	4	0.8%	Yes
—	—	5	1.0%	—
—	—	6	1.2%	—
—	—	7	1.4%	—

TABLE III
GLOBAL UPSCALE EXAMPLE ($E = 8$); ϵ_{step} IS UPDATED FROM 0.1% TO 0.2%.

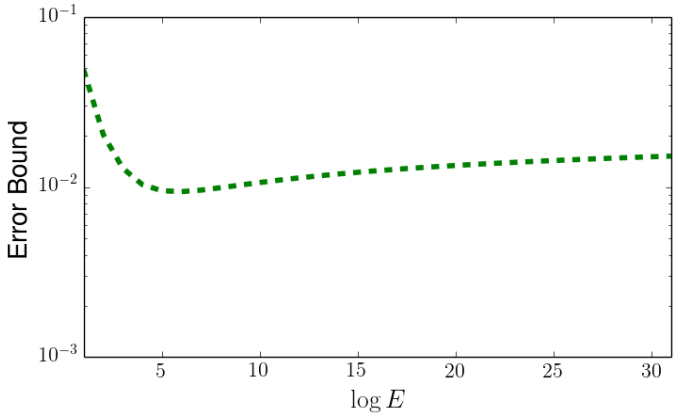

Fig. 6. The error bound from Theorem 6 for N, M of NZ09, different choices of E and an average of 12.5 bits per counter.

E. Parameter Choice

We now describe the process of choosing the parameters to minimize the upper bound from Theorem 5 and then Theorem 6. In the standard scenario, we have limited space for our data structure of T bits. We usually have an upper bound for M, e.g., by multiplying the maximum supported traffic rate by the maximum measurement time. If M is still unknown, we can use the maximum integer we can represent. N could also be given, as the maximum number of flows the networking device supports. If N grows during run-time and we have enough space, we can always allocate more counters on the fly. We now choose L, B, S and E. To minimize ϵ_{max}, we should allocate as many bits as possible for every counter. We therefore allocate $\log_2 L = \lfloor \frac{T}{N} \rfloor$ bits per counter. We are left with $T \mod N$ bits for the scale parameters (if no memory is left we can use a single bucket). Next, we note that there is no point in choosing E to be larger than M. Every upscale should increase the counting capacity by at-least one and therefore M upscales should be always sufficient to achieve the maximum counting capacity. To find the optimal E, we can iterate over $\log_2 E$, which should be an integer number as it represents the number of bits we give the scale parameter. For each choice of E, we calculate the number of buckets we can afford: $B = \lfloor \frac{T \mod N}{\log_2 E} \rfloor$. Given all the parameters, we can calculate the upper bound from Theorem 6. The upper bound requires the computation of the function $\varepsilon(M)$. This function can be computed through binary search because the opposite function is increasing and can be easily computed. By iterating over the possible E values, we can find the E that gives the smallest upper bound. After finding E, we can calculate B and $S = \frac{N}{B}$.

For example, consider the trace NZ09, which will be presented in Section IV. The trace has $N = 32,737,760$ flows. In this example, we allocate 12.5 bits for each counter. M is unknown in advance so we choose the maximum int $2^{32} - 1$. After allocating the maximum of 12 bits per symbol, we are left with $T \mod N = 16,368,880$ bits. We now try all possible options of E, and for each option we calculate B and the upper bound from Theorem 6. Figure 6 depicts the computed bounds for every choice of E. We can see that for Es that are too small, the granularity of the error ϵ_{step} is too crude and the result is a high error. The optimal E for the upper bound is 2^6, and larger Es give higher errors because they require allocation of bigger buckets.

IV. SIMULATION RESULTS

We evaluated ICE-Buckets with five different Internet packet traces. The first trace (NZ09) consists of twenty four hours of Internet traffic collected from an unnamed New Zealand ISP on Jan. 6th 2009 [34]. It is a relatively large trace, containing almost a billion packets and over 32 Million flows. We also used data from two Equinix data-centers in the USA, which are connected to backbone links of Tier1 ISPs. Both use per flow data for load balancing. One is connected to a link between Chicago, IL and Seattle, WA. We used three traces from this dataset, CHI08 from 2008 [35], CHI14 from 2014 [36] and CHI15 from 2015 [37]. CHI08 was previously used to evaluate CEDAR in [31]. Trace (SJ13) was taken in 2013 from an Internet data collection monitor that is connected to a link between San Jose and Los Angeles, CA [38].

We compare ICE-Buckets to two state of the art counter estimation algorithms - DISCO [29] and CEDAR [31]. In order to measure different memory constraints, we tested each algorithm with both 8.5 and 12.5 bits on average per counter. We use 8 and 12 bits (correspondingly) for the symbols. The remainder is used for the scale parameters in ICE-Buckets and for storing the estimation value array in CEDAR. Previous works were evaluated with similar symbol lengths. DISCO does not have an upscaling scheme. Therefore we configured it according to the maximal expected number of packets (M) that is different for each trace, as specified in Table IV. Per-trace statistics and configurations are given in Table IV.

Table V presents the overall relative error of ICE-Buckets and the alternatives for the tested traces. We also present the upper bounds on the error of ICE-Buckets and CEDAR. As can be observed, for real datasets, ICE-Buckets' error is much lower than this bound, since the majority of flows are small. For CHI08, ICE-Buckets achieves an overall relative error that is over 57 times smaller than that of CEDAR. Notice that for all traces, ICE-Buckets' overall relative error with 8-bit symbols is lower than that of the alternatives, even with 12-bit symbols. Note that in our case, DISCO is 100% accurate when the value is 1, and is slightly less accurate than CEDAR

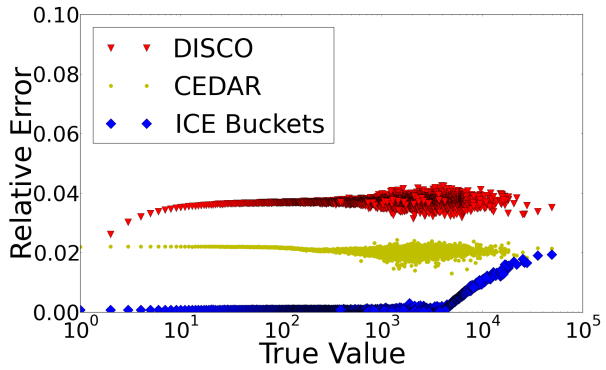

(a) CHI08 with 12-bit symbols

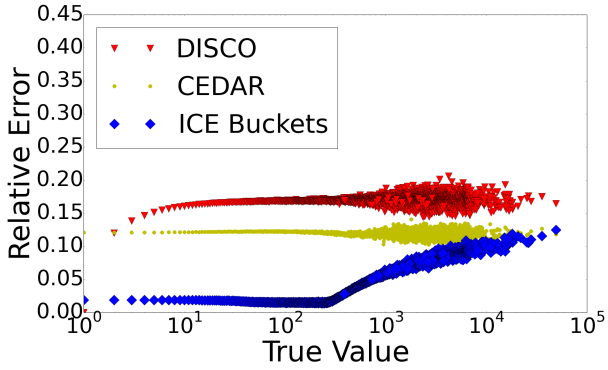

(b) CHI08 with 8-bit symbols

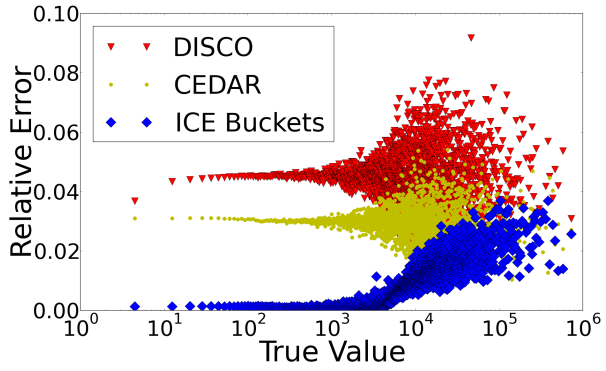

(c) NZ09 with 12-bit symbols

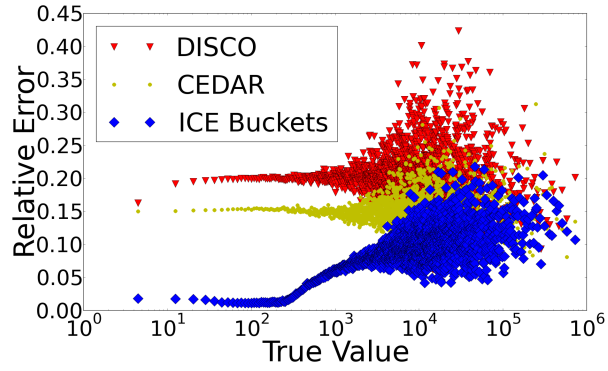

(d) NZ09 with 8-bit symbols

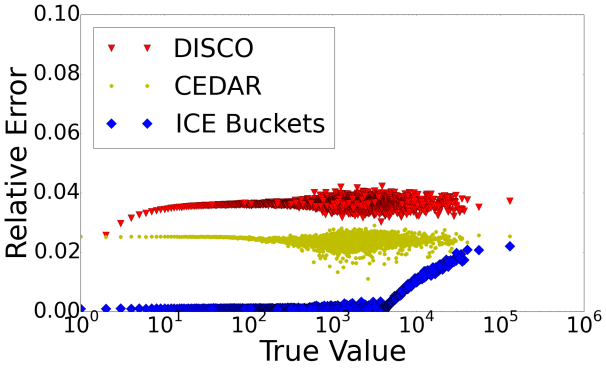

(e) CHI15 with 12-bit symbols

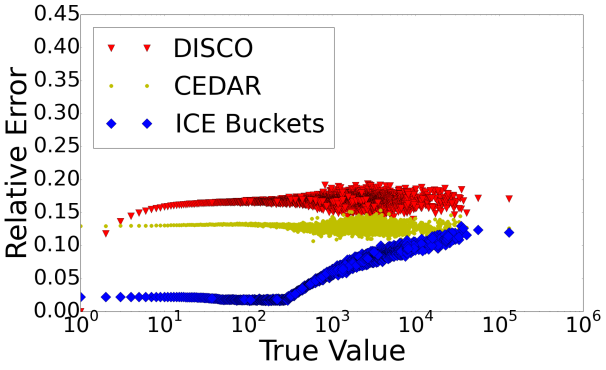

(f) CHI15 with 8-bit symbols

Fig. 7. Comparison of the relative error per value on different traces using 8 and 12-bit symbols

for all other values. All in all, this results in an overall relative error similar to that of CEDAR.

We also experimented with a version of ICE-Buckets that does not use global upscale, which should be easier to implement in hardware. To do so, we pre-configured ϵ_{step} to ensure local upscales are sufficient to count 2^{32}-1 packets in any bucket. Note that for most traces the error in this case is very similar to that of ICE-Buckets with upscale. We therefore recommend to implement ICE-Buckets without global upscale when the total number of packets can be bounded in advance.

To explain the cause of ICE-Buckets' substantial error reduction, we show in Figure 7 the relative error as a function of the real counter value. The relative error was computed from 256 runs of each algorithm on CHI08 and CHI15 and one run on NZ09. Note that for most counter values, the relative error of ICE-Buckets is lowest, followed by CEDAR, and then DISCO. ICE-Buckets achieves an error close to zero for counters smaller than L because an accurate counter of $log_2 L$ bits suffices to represent those values. Unfortunately, this error cannot always be zero, as some of these counters share buckets with larger counters. As the counter scale grows, the estimation error increases, until eventually, the largest counter is estimated with ϵ_{max}. In contrast, CEDAR estimates all of the counters with relative error ϵ_{max}.

Trace	Flows (N)	Packets	M	S 8b / 12b	E 8b / 12b
CHI08	1,420,318	26,750,712	26,750,712	10 / 14	32 / 128
CHI14	1,213,614	34,721,808	34,721,808	10 / 12	32 / 64
SJ13	3,071,187	20,803,060	20,803,060	12 / 14	64 / 128
NZ09	32,737,760	891,023,765	$2^{32}-1$	12 / 12	64 / 64
CHI15	683,708	18,774,214	18,774,214	10 / 12	32 / 64

TABLE IV
DOCUMENTATION OF TRACE CHARACTERISTICS AND ICE-BUCKETS CONFIGURATIONS USED IN EXPERIMENTS FOR 8/12 BIT SYMBOLS.

Trace	CHI08		CHI14		SJ13		NZ09		CHI15	
Bits-Per-Symbol	8	12	8	12	8	12	8	12	8	12
CEDAR (upper bound)	16.93%	3.70%	17.10%	3.75%	16.76%	3.65%	19.61%	4.51%	16.70%	3.64%
ICE-Buckets (upper bound)	5.02%	0.42%	5.71%	0.50%	3.49%	0.26%	8.22%	0.94%	5.63%	0.49%
DISCO (actual)	10.34%	2.24%	10.16%	2.24%	7.82%	1.71%	14.24%	3.21%	11.58%	2.51%
CEDAR (actual)	12.19%	2.17%	12.40%	2.38%	13.05%	2.60%	15.01%	3.09%	12.96%	2.51%
ICE-Buckets (no global upscale)	1.70%	0.10%	2.14%	0.14%	1.00%	0.03%	1.78%	0.14%	4.29%	0.21%
ICE-Buckets (actual)	1.50%	0.06%	1.96%	0.16%	1.01%	0.03%	1.81%	0.13%	2.02%	0.11%

TABLE V
OVERALL RELATIVE ERROR BOUNDS AND ACTUAL OVERALL RELATIVE ERROR OF DIFFERENT ALGORITHMS ON VARIOUS TRACES

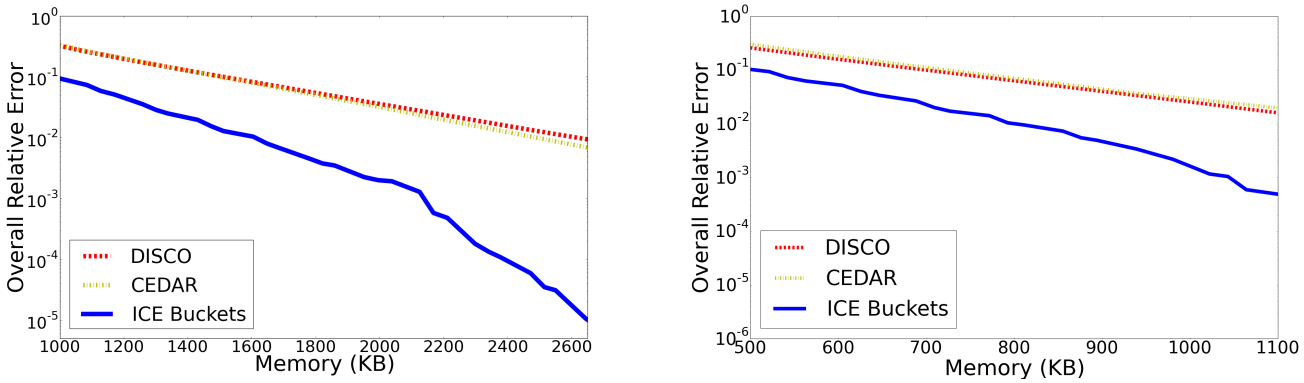

(a) CHI08 (b) CHI15

Fig. 8. A comparative evaluation of the accuracy as a function of the allocated memory space; overheads are taken into account!

Figure 8 illustrates the accuracy of different algorithms for CHI08 and CHI15 under varying memory constraints. Overheads of all methods are taken into account and the maximal symbol size is used for each method. ICE-Buckets uses different configurations with overheads that range between $\frac{1}{4}$ and $\frac{3}{4}$ bits per counter. The figure shows that DISCO and CEDAR provide similar space-accuracy trade-offs as mentioned in Section II-D3. Under all of the simulated memory constraints, ICE-Buckets is more accurate than both CEDAR and DISCO, and the difference between ICE-Buckets and the other algorithms grows with the memory. We explain this by noting that as the number of bits per symbol (L) grows, more counters can be estimated with zero error.

Figure 9 describes the overall relative error of CEDAR and ICE-Buckets throughout the NZ09 trace's progress. To adapt to the growing counter scale, both ICE-Buckets and CEDAR use an upscale mechanism that gradually increases the error. Note that the overall relative error of ICE-Buckets is almost constant throughout an entire day of real Internet traffic. In addition, CEDAR's multiple global upscales are clearly visible in the figure. In contrast, ICE-Buckets' upscales are mostly local and cause a smoother increase in the relative error.

V. CONCLUSION

In this work, we have introduced ICE-Buckets, a novel counter estimation data structure that minimizes the relative error. ICE-Buckets uses the optimal estimation function with a scale that is optimized independently for each bucket.

We first described an explicit representation of this function, which was previously known only in recursive form. We extended its analysis and showed a method to measure the effect of upscale operations on the relative error. This function is used in ICE-Buckets to minimize the error in each bucket.

ICE-Buckets is dynamically configured to adapt to the growing counters. For practical deployments, it can be implemented without global operations while providing similar accuracy.

We proved an upper bound to ICE-Buckets' overall relative error, which is significantly smaller than that of previous estimation algorithms. In particular, we demonstrated a reduction

of up to 14 times in this upper bound when applied to traffic characteristics of real workloads. ICE-Buckets also achieves the same maximum relative error as the optimal function.

Additionally, we extensively evaluated ICE-Buckets with four Internet packet traces and demonstrated a reduction of up to 57 times in overall error. ICE-Buckets achieves an improvement in accuracy even when it is given considerably less space than the alternatives. Finally, we have shown that ICE-Buckets is significantly more accurate than the leading alternatives for a wide range of memory constraints.

In this work, we explained how to perform decrements and downscaling. Yet, doing so, greatly complicates the analysis. Analyzing their impact is left for future work.

As mentioned before, another interesting topic for future work is combining shared counters schemes like CMS [9], multi-stage filters [24], SBF [25], as well as TinyTable [26] with estimators. Since ICE-buckets offers small counters with low error, replacing the counters in the above with estimators could potentially improve their space to accuracy ratio. Another promising aspect of the above is that estimators can return the estimated value in $O(1)$ time, thereby maintaining the access efficiency of such combined schemes.

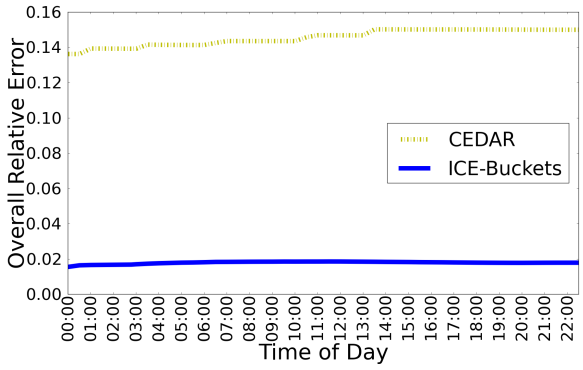

Fig. 9. Average relative error through trace progress with 8-bit symbols as simulated on NZ09

Acknowledgements: We would like to thank Isaac Keslassy, Olivier Marin and Erez Tsidon for their insights and helpful advice. This work was partially funded by the Israeli Ministry of Science and Technology grant 3-10886 and the Technion HPI center.

REFERENCES

[1] D. Shah, S. Iyer, B. Prabhakar, and N. McKeown, "Maintaining statistics counters in router line cards." *IEEE Micro*, vol. 22, no. 1, pp. 76–81, 2002.

[2] S. Ramabhadran and G. Varghese, "Efficient implementation of a statistics counter architecture," *ACM SIGMETRICS*, vol. 31, no. 1, pp. 261–271, 2003.

[3] A. Kabbani, M. Alizadeh, M. Yasuda, R. Pan, and B. Prabhakar, "Af-qcn: Approximate fairness with quantized congestion notification for multi-tenanted data centers," in *Proc. of the 18th IEEE Symposium on High Performance Interconnects*, ser. HOTI, 2010, pp. 58–65.

[4] B. Mukherjee, L. Heberlein, and K. Levitt, "Network intrusion detection," *IEEE Network*, vol. 8, no. 3, pp. 26–41, 1994.

[5] P. Garcia-Teodoro, J. E. Daz-Verdejo, G. Maci-Fernndez, and E. Vzquez, "Anomaly-based network intrusion detection: Techniques, systems and challenges," *Computers and Security*, pp. 18–28, 2009.

[6] G. Dittmann and A. Herkersdorf, "Network processor load balancing for high-speed links," in *Proc. of the Int. Symposium on Performance Evaluation of Computer and Telecommunication Systems*, 2002.

[7] G. Einziger and R. Friedman, "Tinylfu: A highly efficient cache admission policy," in *Euromicro PDP*, 2014, pp. 146–153.

[8] C. Estan and G. Varghese, "New directions in traffic measurement and accounting," in *Proc. of the ACM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, ser. SIGCOMM, 2002, pp. 323–336.

[9] G. Cormode and S. Muthukrishnan, "An improved data stream summary: The count-min sketch and its applications," *J. Algorithms*, vol. 55, pp. 29–38, 2004.

[10] A. Ciuffoletti, "Architecture of a network monitoring element," FORTH-ICS, Tech. Rep., 2006.

[11] V. Paxson, J. Mahdavi, A. Adams, and M. Mathis, "An architecture for large scale internet measurement," *IEEE Communications Magazine*, vol. 36, no. 8, pp. 48–54, 1998.

[12] C. Estan, K. Keys, D. Moore, and G. Varghese, "Building a better netflow," in *ACM SIGCOMM*, 2004, pp. 245–256.

[13] Y. Lee and Y. Lee, "Toward scalable internet traffic measurement and analysis with hadoop," *ACM SIGCOMM*, vol. 43, no. 1, pp. 5–13, Jan. 2012.

[14] H. Wang, H. Zhao, B. Lin, and J. Xu, "Dram-based statistics counter array architecture with performance guarantee," *IEEE/ACM Transactions on Networking*, vol. 20, no. 4, pp. 1040–1053, Aug 2012.

[15] Y. Lu, A. Montanari, B. Prabhakar, S. Dharmapurikar, and A. Kabbani, "Counter braids: a novel counter architecture for per-flow measurement," in *ACM SIGMETRICS*, 2008, pp. 121–132.

[16] T. Li, S. Chen, and Y. Ling, "Per-flow traffic measurement through randomized counter sharing," *IEEE/ACM Transactions on Networking*, vol. 20, no. 5, pp. 1622–1634, Oct 2012.

[17] M. Chen and S. Chen, "Counter tree: A scalable counter architecture for per-flow traffic measurement," in *Proc. of IEEE International Conference on Network Protocols (ICNP)*, 2015.

[18] N. Hua, B. Lin, J. J. Xu, and H. C. Zhao, "Brick: A novel exact active statistics counter architecture," in *Proc. of the ACM/IEEE ANCS*, 2008, pp. 89–98.

[19] B.-Y. Choi, J. Park, and Z.-L. Zhang, "Adaptive random sampling for load change detection," *ACM SIGMETRICS*, pp. 272–273, 2002.

[20] Y. Zhang, S. Singh, S. Sen, N. Duffield, and C. Lund, "Online identification of hierarchical heavy hitters: Algorithms, evaluation, and applications," in *Proc. of the ACM IMC*, 2004, pp. 101–114.

[21] A. Metwally, D. Agrawal, and A. E. Abbadi, "Efficient computation of frequent and top-k elements in data streams," in *Proc. of the 10th International Conference on Database Theory (ICDT)*, 2005.

[22] R. M. Karp, S. Shenker, and C. H. Papadimitriou, "A simple algorithm for finding frequent elements in streams and bags," *ACM Trans. Database Syst.*, vol. 28, no. 1, Mar. 2003.

[23] R. Ben-Basat, G. Einziger, R. Friedman, and Y. Kassner, "Heavy hitters in streams and sliding windows," in *Proc. of the IEEE Conference on Computer Communications (INFOCOM)*, 2016.

[24] D. Ficara, A. Di Pietro, S. Giordano, G. Procissi, and F. Vitucci, "Enhancing counting bloom filters through huffman-coded multilayer structures," *Networking, IEEE/ACM Transactions on*, vol. 18, no. 6, pp. 1977–1987, 2010.

[25] S. Cohen and Y. Matias, "Spectral bloom filters," in *ACM SIGMOD*, 2003, pp. 241–252.

[26] G. Einziger and R. Friedman, "Counting with TinyTable: Every Bit Counts!" in *Proc. of the International Conference on Distributed Computing and Networking*, ser. ICDCN, 2016.

[27] R. Morris, "Counting large numbers of events in small registers," *Commun. ACM*, vol. 21, no. 10, pp. 840–842, 1978.

[28] R. Stanojevic, "Small active counters," in *IEEE INFOCOM*, 2007, pp. 2153–2161.

[29] C. Hu, B. Liu, H. Zhao, K. Chen, Y. Chen, Y. Cheng, and H. Wu, "Discount counting for fast flow statistics on flow size and flow volume," *IEEE/ACM Transactions on Networking (TON)*, vol. 22, no. 3, pp. 970–981, 2014.

[30] C. Hu and B. Liu, "Self-tuning the parameter of adaptive non-linear sampling method for flow statistics," in *Proc. of Int. Conf. on Computational Science and Engineering (CSE)*, Aug 2009, pp. 16–21.

[31] E. Tsidon, I. Hanniel, and I. Keslassy, "Estimators also need shared values to grow together," in *IEEE INFOCOM*, 2012, pp. 1889–1897.

[32] L. Yang, W. Hao, P. Tian, D. Huichen, L. Jianyuan, and L. Bin, "Case: Cache-assisted stretchable estimator for high speed per-flow measurement," in *IEEE INFOCOM*, 2016, pp. 26–30.

[33] C. Hu, S. Wang, B. L. Tian, Jia, Y. Cheng, and Y. Chen, "Accurate and efficient traffic monitoring using adaptive non-linear sampling method," in *IEEE INFOCOM*, 2008, pp. 26–30.
[34] W. N. R. Group, "Waikato Internet Traffic Storage 2009-01-06 00:00-23:30 UTC, ISPDSL-I," http://wand.net.nz/wits/ispdsl/1/.
[35] P. Hick, "CAIDA Anonymized 2008 Internet Trace, equinix-chicago 2008-03-19 19:00-20:00 UTC, Direction A."
[36] ——, "CAIDA Anonymized 2014 Internet Trace, equinix-chicago 2014-03-20 13:55 UTC, Direction B."
[37] ——, "The CAIDA UCSD Anonymized Internet Traces - 2015-02-19 13:00 UTC, Direction A."
[38] ——, "CAIDA Anonymized 2013 Internet Trace, equinix-sanjose 2013-1-17 13:55 UTC, Direction B."